



Multi-Camera Persistent Surveillance Test Bed

**by David Baran, Barry O'Brien, Nick Fung,
Jesse Kovach, and David Miller**

ARL-TR-4031

January 2007

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Adelphi, MD 20783-1197

ARL-TR-4031

January 2007

Multi-Camera Persistent Surveillance Test Bed

**David Baran, Barry O'Brien, Nick Fung,
Jesse Kovach, and David Miller
Computational and Information Sciences Directorate, ARL**

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) January 2007		2. REPORT TYPE Final		3. DATES COVERED (From - To) March 2006 to November 2006	
4. TITLE AND SUBTITLE Multi-Camera Persistent Surveillance Test Bed				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) David Baran, Barry O'Brien, Nick Fung, Jesse Kovach, and David Miller				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory ATTN: AMSRD-ARL-CI-CB 2800 Powder Mill Road Adelphi, MD 20783-1197				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-4031	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory 2800 Powder Mill Road Adelphi, MD 20783-1197				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT <p>Recognizing the U.S. military's superiority in open battlefield environments, adversaries have moved the battle into cities, generating a different set of technical challenges for the modern war fighter. Combat in urban environments characterized by large civilian populations and high building densities requires a different tactical approach to ground operations. Radical new solutions are required to reduce the number of dangerous situations our Soldiers encounter by providing them with improved situational awareness. Small reconnaissance, surveillance, and target acquisition (RSTA) platforms have the ability to cooperate through information sharing to increase the situational awareness over a region of interest (ROI). One RSTA task that can be performed by such platforms is persistent surveillance - the ability to monitor objects of interest without interruption over a large ROI.</p> <p>This paper discusses the design and implementation of a persistent surveillance test bed comprised of a homogenous group of stationary assets. Further, it examines persistent surveillance algorithms that can be reduced to practice, and their proposed implementation and evaluation within the test bed.</p>					
15. SUBJECT TERMS Persistent surveillance, image processing, RSTA					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 30	19a. NAME OF RESPONSIBLE PERSON David Baran
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include area code) (301) 394-5613

Contents

List of Figures	iv
Introduction	1
System Overview	2
Camera and Enclosure	2
Mounting of Camera Enclosure	3
Conversion Box Design	4
System Integration	8
Components	9
Software Overview	10
Existing Software	10
New and Modified Software	12
Surveillance Applications	14
Tracking Applications	18
Heterogeneous Asset Collaboration	19
Conclusion and Future Work	20
Distribution List	23

List of Figures

Figure 1. Mounted camera enclosure.....	4
Figure 2. Internal configuration of conversion box.	6
Figure 3. Bulkhead connectors on exterior of conversion box.	7
Figure 4. Internal configuration of re-designed conversion box.....	8
Figure 5. DVR client application.	14
Figure 6. Camera GAC showing an obstruction.....	15
Figure 7. Coverage area of two cameras.....	16
Figure 8. Elevation view of GAC.	17

Introduction

Urban warfare poses a different set of challenges for the U.S. military. Recognizing the U.S. military superiority in open battlefield environments, adversaries have moved the battle into cities, generating a different set of technical challenges for the modern war fighter. Combat in urban environments characterized by large civilian populations and high building densities requires a different tactical approach to ground operations. Traditional technologies and techniques fail in these constrained environments. Radical new solutions are required to reduce the number of dangerous situations our Soldiers encounter by providing them with improved situational awareness.

To work effectively in this complex urban landscape, the Army is reorganizing into smaller, more agile and more independent combat units. To complement this regrouping, reconnaissance and surveillance systems need to be similarly reorganized into small, agile, independent working units that can operate in an autonomous or semiautonomous fashion to create and maintain good situational awareness.

Teams of small RSTA sensor platforms can cooperate to perform tasks more effectively than a single platform in isolation. A greater amount of information about the environment can be obtained with reduced uncertainty through cooperation and the sharing of situational awareness knowledge obtained from the fusion of multiple viewpoints. Furthermore, a system consisting of a set of distributed sensor platforms is capable of observing and interacting with a larger region of the battle space, is resilient to failures of individual components, and is able to quickly adapt to a changing environment.

Persistent surveillance is one of several RSTA tasks that can be performed by such a system. Various types of assets are required in order to provide adequate persistent surveillance for small combat units operating within complex terrain over a large region of interest (ROI). An asset is any capability that aids or enhances the ability of a combat unit to perform its mission, including humans, software, and hardware capabilities such as sensor platforms. Assets contain one or more basic capabilities that, when integrated, provide a more complex capability that performs a useful function as a higher level asset. When tasked, higher level assets can provide domain and/or situation specific functionality under the direction and guidance of a leader. An asset may have a dedicated platform or may exist on some other asset's platform. As such, it may or may not have mobility. Assets have inherent built-in behaviors that enable them to exist and perform useful tasks under the direction and guidance of a centralized leader or via distributed control. An asset may have the ability to discover and utilize information and/or functionality from other assets.

In order to effectively evaluate and implement persistent surveillance algorithms, a test bed consisting of a small number of identical assets was required. This paper discusses the design and implementation of a persistent surveillance test bed comprised of a homogenous group of stationary assets. Further, it examines persistent surveillance algorithms that can be reduced to practice, and their proposed implementation and evaluation within the test bed.

System Overview

The prototype persistent surveillance test bed is installed on the roof of the main laboratory buildings of the Army Research Laboratory (ARL) at Adelphi Laboratory Center (ALC). This location was chosen as it complements other existing surveillance assets while providing a convenient testing location. In order to support surveillance capability on all sides of the building, four cameras were used, one positioned at each corner of the building roof. Due to their distributed locations, power and fiber cables were run to each of the cameras. Extension cords to reach between existing AC power outlets, which were used for the system, and the desired camera locations were custom made, using 16 AWG¹ 3-conductor neoprene jacketed power cable (Belden #19208). The wall outlet end of the cable was terminated with a male Woodhead connector (Woodhead #14W47-Blk), and the camera end of the cable was terminated with a female Hubbell connector (Hubbell #HBL52CM69C). Fiber-optic cable was run through an existing conduit from a laboratory bay to the roof, and then to the four camera locations. The custom length multimode fiber cables had an OFNR² jacket, and were selected to provide a weatherized solution.

Camera and Enclosure

A central part of a persistent surveillance system is a capable imaging device. The Battlefield Information Processing Branch (CI-CB) of ARL has previously used the DI-5000 camera from ICx Digital Infrared Imaging, Inc. with some success on both mobile and stationary platforms. The DI-5000 is a capable device, including both color and thermal imagers mounted on an unconstrained pan-tilt unit. However, experience has shown that the cameras have been damaged under poor weather conditions, and their high initial cost (\$15K-\$20K) has proven to be prohibitive. The DI-5000 also lacks a built-in video digitizer and server, producing only NTSC analog video.

¹American Wire Gauge.

²Optical Fiber Nonconductive Riser.

The camera system selected for this application was the Sony Network Camera (Sony #SNC-R230N) commonly referred to as a Sony Ball-Cam. The Sony Ball-Cam includes color and low-light near-infrared imaging capabilities on a 340 degree pan and 115 degree tilt gimbal. Also featured on the Sony Ball-Cam is embedded Ethernet support, including a built-in web interface, simplifying integration of the camera system.

This system was intended to be deployed for extended periods of time in all weather conditions and the Sony Ball-Cam is not weather-resistant. Consequently, the camera was installed in a protective enclosure. Dotworkz Systems produces a climate-controlled enclosure called the Cooldome³ that is designed to house a Sony Ball-Cam. In addition to protecting the camera from precipitation, the Cooldome contains a thermo-electric cooling system, called a Peltier cooler, that cools the enclosure without exchanging air with the environment. Peltier coolers are semiconductor devices that utilize the Peltier Effect, a phenomenon that causes electrons to speed up or slow down under the influence of a contact potential difference. For example, if you put a drop of water in the hollow on the joint of two semiconductors and run current through the material, the drop will freeze; if the direction of the current is reversed the droplet will vaporize⁴. The temperature change created and ultimately the effectiveness of a Peltier cooler depends on the kinetic energy of electrons. When electrons slow down, their kinetic energy decreases, thus cooling the air between the semiconductors. By monitoring the ambient temperature, the cooler can be used to keep the temperature of the camera within an acceptable range. Using this method of cooling, the Cooldome is able to protect the enclosed camera in ambient temperatures up to 145° F, and can create up to a 45° F drop between the ambient and internal temperatures.

While the use of the Sony Ball-Cam/Dotworkz system sacrifices true thermal imaging capabilities and some performance, its weather resistance and low cost (less than \$3K) made it an acceptable choice for this prototype system.

Mounting of Camera Enclosure

CI-CB has previously mounted a DI-7000⁵ on the roof to act as a surrogate aerostat vehicle payload. Using a custom mounting bracket, we attached the camera to a long, horizontal beam that extended outward from the side of the building, allowing the camera to survey a greater amount of area. While fully functional, the weight and profile of the camera caused it to be susceptible to the effects of wind. This caused the image from the camera to shake.

The Dotworkz enclosure was originally intended to mount directly to the side of a building, not suspended over the edge of one, so modifications were required. Previous experience

³Cooldome is a trademark of Dotworkz System.

⁴<http://www.digit-life.com/articles/peltiercoolers>.

⁵The DI-7000 is a newer version of the DI-5000 camera.

demonstrated that the camera should be mounted on a vertical beam as close to the edge of the roof as possible, as shown in figure 1. This left only one acceptable mounting location for the beam on the roof: the railing that encompasses the entire roof top. The beam was attached to the railing using two large U-bolts, allowing it to be fastened securely and prevent wind-born movement. Although there was initially concern about the vertical field of view being obscured by mounting the camera near the edge of the building this turned out to not be an issue. The closer position afforded a complete view of the ground adjacent to the building directly below the camera.



Figure 1. Mounted camera enclosure.

Conversion Box Design

Once the cameras were mounted on the roof, the necessary power and data connections had to be constructed and connected to the cameras. As previously described, AC power and Ethernet connections via fiber optic cable were installed to each camera location. In order for the camera systems to use these connections, several conversion devices had to be incorporated.

The Sony Ball-Cam/Dotworkz system operated on 12V DC power, requiring the standard 110V AC power provided from the wall outlet to be converted. Based on the original power requirements determined for the camera system, we selected a Kepco AC to DC converter

(Kepco #FAW 12-4.2K) that output 50W max at 4.2A. We chose this model because its output amperage was approximately twice the required amount, allowing the converter to operate cooler.

Similarly, the system required an Ethernet connection in order to transfer data and incorporate the individual cameras into the larger system. Ideally, Cat.5 cable would have been used, as the camera has a built-in 100 Base T Ethernet interface. However, the distance between the camera locations and server (up to 255m) was too large for Cat.5 cable to effectively support⁶. We routed fiber optic cable to the cameras instead of Cat.5 Ethernet cable, as fiber is not susceptible to the same path loss rates as Cat.5. Handling the fiber optic to Cat.5 conversion was the Unicom Velocity Media Converter⁷, selected for its low cost.

Since the Kepco and Unicom converters were not weatherproof, a protective enclosure was required. Instead of investigating a costly, custom weatherproof enclosure, a Pelican case (Pelican Case #1400) was used. Pelican cases are known for their extreme ruggedness and weather resistance, and are available in a variety of shapes and sizes. The box had to enclose the two conversion devices. Thus, the size of the box was based on their dimensions. Additionally, the amount of free air space that would be present inside of the box was considered. Free space inside of the box would help dissipate some heat generated by the power converter.

The Unicom media converter required 5V DC power. Since it also came equipped with an AC adapter wall plug, we installed an AC receptacle in the box. This receptacle negated the need to add an additional component to the box's power system in order to convert the 12V output of the power converter to 5V. Instead, the receptacle only required simple wiring from the input of the AC to DC converter while occupying a minimal amount of space in the enclosure.

In order to secure the components to the case, we designed a metal mounting plate. The plate dimensions mirrored those of the Pelican case, with one corner being cut out to allow for easy removal of the plate from the box. Each of the components was screw-mounted onto the plate, and the plate was attached to the bottom of the box with heavy-duty Velcro⁸ strips as shown in figure 2. This plate increased the ruggedness of the system, by restricting the the components ability to move inside of the case during system transport.

After the design of the conversion box interior was completed, a method to get the required cables into and out of the box while still retaining its weather resistance and portability needed to be designed. The box required four connections: AC power in from the wall, DC power out to the camera, fiber optic lines to communicate with the central server, and Cat 5 Ethernet cable to

⁶Cat.5 cable has a supported maximum distance of approximately 100 m.

⁷Velocity is a trademark of Unicom Electric, Inc.

⁸Velcro is a registered trademark of Velcro USA, Inc.




Figure 2. Internal configuration of conversion box.

communicate with the camera. Not only was it important to keep the components inside of the case dry, but the connections for these cables also had to be weatherized to avoid shorting out the electrical signals.

For the DC power interface, we used a Marinco receptacle/plug system (#12VPK) intended for marine applications. The Marinco interface consisted of a positive locking plug with a twist-disconnect mechanism. The AC power interface used a waterproof receptacle and associated plug from Hubbell (#HBL61CM64 and #HBL52CM69C). We purchased a waterproof boot also made by Hubbell that surrounds the plug to make a water-resistant seal when a male connector is inserted into the female plug, but this boot had too large of a diameter to work with the plug-receptacle pair. However, system testing demonstrated that the seal between the plug and receptacle was tight enough to protect the system, except for an extreme case of a prolonged, wind-driven rain event.

For the Ethernet connection, we used a weather resistant Cat.5 quick disconnect connector, the Woodhead Cat.5E Industrial Shielded Feed-Thru Coupler and Plug Kit (#ENSP1F5 and #ENSAM315). A similar weather resistant quick disconnect for the fiber optic cable was not commercially available. Instead, an Olflex pass thru connector (#PG-11) was used. When tightened, the Olflex connector created a reasonable seal using a neoprene washer around the fiber optic cable. While this seal worked well in protecting the inside of the box, it was very

difficult to insert or remove the fiber optic cables from the connector. All of the connectors are shown in figure 3.



Figure 3. Bulkhead connectors on exterior of conversion box.

After the initial installation of the cameras and conversion boxes, it became apparent that the design could not supply the correct amount of current needed by the Peltier cooler during operation. Instead of the approximately 2A the original system was designed to provide, we determined that when the Peltier cooler was activated to cool the camera enclosure the system required almost 13A. The Peltier coolers were temporarily disabled until revisions could be made to the conversion box to support the required power output.

To correct the power issue, several components of the system had to be redesigned. A new Kepco AC to DC power converter (Kepco #RAX 12-14K) capable of delivering 15.4 amps at 175W replaced the original system power converter. Since the new power converter was physically much larger than the original, a larger Pelican case (Pelican #1520) was provided to accommodate the increased converter size. Once again free space was taken into account to ensure proper heat dissipation from the power supply inside the case. We designed a new mounting plate to fit the case, in a similar fashion as previously described. The new case and mounting plate are shown in figure 4.




Figure 4. Internal configuration of re-designed conversion box.

Along with a larger capacity power converter, the DC plug and receptacle had to be replaced with a higher rated version to handle the increased amperage. While the original connectors were rated to 15A, they contained a 10A fuse. Once again, we chose marine grade components with a twist disconnect mechanism in order to provide a weatherized power solution. The new DC connector (Marinco #540-0005) and plug (Marinco #540-0006) were rated to 30 amps.

While performing system upgrades, other minor issues were addressed. One of these was the O-flex connector used for the fiber cables. The original O-flex connector was replaced with a version one size larger (Olflex #PG-13). This still provided a weatherized, watertight fit when secured but allowed for easier system set-up.

System Integration

The final surveillance system test bed consisted of the four roof-top camera assets and connecting infrastructure, one server responsible for video processing and recording, a number of user workstations, and connecting infrastructure. The four cameras are connected to an internet protocol (IP) network using fiber optic cable and fiber/100BaseT Ethernet media converters, along with the video processing server and the user workstations. All data, including live and recorded digitized video and control traffic, is distributed via the IP network. Future automated

surveillance applications and processing components will also communicate with the camera infrastructure through this IP network.

Components

All video processing for the camera system is handled by a standard Windows⁹-based server; an IBM xSeries¹⁰ 336 equipped with a 3.6 GHz Intel Xeon¹¹ processor, 1 gigabyte (GB) of Random Access Memory (RAM), and 110 GB of high-speed small computer system interface (SCSI) storage. Software running on this computer reads data from the four cameras via hypertext transfer protocol (HTTP) and makes it available using ARL-developed protocols suitable for this surveillance application. The server can perform image processing operations including stabilization and target tracking on the video streams. The server software also supports archiving and retrieval of video streams from the roof cameras along with other surveillance assets.

Operator stations are standard Windows-based computers connected to the network running client software developed for the surveillance system. There can be any number of operator stations connected to the system within the resource limits of the network. Additionally, the operator stations can be equipped with geographic information system (GIS) mapping software to support advanced automated surveillance and situation display applications.

An extensible architecture allows new components to be easily integrated into the system. Such components range from new image processing algorithms and additional types of video sources to components that automatically control the cameras and otherwise assist the operator.

All video sent over the network is compressed in Motion-JPEG¹² format. This video compression format was used for two reasons: it is the native compression format supported by the Sony network cameras and each frame is represented as a discrete JPEG image. Since each frame has no dependencies on previous or subsequent frames, storage and retrieval of individual frames at a point in time is simplified. Additionally, motion-JPEG has advantages when used on unreliable networks since the lack of interdependencies between frames minimizes the effect of packet loss.

⁹Windows is a trademark of Microsoft Corporation.

¹⁰IBM xSeries is a registered trademark of IBM Corporation.

¹¹Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation.

¹²Joint Photographic Experts Group.

Software Overview

The software used in the persistent surveillance test bed consists mainly of components and applications developed at ARL. Much of this software was pre-existing, having been developed for previous programs such as Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance (C4ISR) On the Move and Horizontal Fusion as well as in-house research projects. These existing software components were used to provide service discovery, network streaming video, video display, and pan/tilt camera control capabilities. Some modifications to the existing software were required for this effort. Additionally, new software components were developed for video stabilization, motion detection and tracking, and video archiving and retrieval. A modular system architecture allowed for the rapid integration of existing code and new components to produce a functional system.

Existing Software

Over the years, CI-CB has developed a modular suite of software originally designed for battlefield command and control (C2) systems that has proven itself to be easily extensible for use in many related areas. Components have been developed for various functions including blue/red force tracking, sensor fusion, geographic situation display and mapping, and tele-operation of robotic vehicles. These components are built on top of an ARL-developed core network library known as CipNet¹³, which provides basic reliable and unreliable message delivery functionality along with a service discovery framework.

The CipNet communication and discovery framework is organized around the concept of platforms consisting of one or more hosts. Hosts are grouped into platforms manually by the system integrator. For example, all of the systems on an unmanned ground vehicle (UGV) would be organized as a single platform. This UGV platform would advertise the services provided by these systems, such as platform mobility, pan-tilt control, video, and posture/position information. Platforms are assumed to occupy a discrete geographic location. Hosts within a platform are assumed to have reliable high-speed connectivity with each other. The CipNet libraries along with the underlying IP stack handle the mapping between the services advertised by a platform and individual programs running on platform hosts.

CipNet currently uses a hierarchical system, known as the agent registry server (or aregserver), to handle service discovery. Each platform runs one instance of the agent registry server. Additionally, one platform designated by the system operator hosts the master agent registry,

¹³Choy, S. "CIP Network Agent Service API User's Guide" ARL, Adelphi, MD.

which maintains the authoritative list of all services available within the system. Servers running on a platform send service advertisements to the platform's agent registry, which forwards any updates to the master agent registry. The master agent registry then updates its list of services and propagates the changes to the agent registries running on other platforms. Platform agent registries use these notifications to maintain a cache of the service list for the entire network. To discover servers, clients query the agent registry server on their platform, which responds with a list of services from the local cache that matches the request. This architecture requires that the address of the host running the master agent registry be statically configured on each platform. All other aspects of configuration, port assignment, and discovery are dynamic.

This discovery model is well-suited to the persistent surveillance system, as a master agent registry server can be easily chosen. Connectivity to the master registry can be guaranteed due to the local nature of the system. Devices, services, and operator stations can be added and removed from the system dynamically without requiring a restart of the entire system. Additionally, the system facilitates experimentation as a given component can be repeatedly stopped and started for debugging, modification, or tuning purposes without affecting the operation of the rest of the system.

The surveillance system also makes use of a preexisting video server application known as MCVideoAgent. MCVideoAgent is a portable C++ application that acquires video from a video source, tags each frame with metadata, and transfers the data stream to clients via the network. The server can also perform some image processing operations on the video frames. MCVideoAgent is built on a modular architecture based around the concept of a “plug-in chain”. The first plug-in in the chain acquires a frame from a video source, and subsequent plug-ins in the chain operate on the frame. Such operations include tagging the frame with posture metadata, processing the image, and compressing the frames when required. The plug-in architecture allows for the rapid addition of new processing plug-ins and video sources.

The surveillance system also uses a preexisting component called CollectControl, a video client and asset control application originally developed for robotic C2 applications. CollectControl is a Windows application written using C# that allows the user to view and record video from any video source on the network. CollectControl queries the agent registry for video sources and presents the user with a list of available cameras. When the user selects a camera from the list, the application begins to display video and allows the user to control any pan/tilt or mobility functions available on the camera platform. Users are able to record video clips or capture and annotate still images. Video clips are saved in standard audio video interleave (AVI) format while stills are saved in a standard JPEG format. Posture data is stored along with all saved video clips and still images through the use of an extensible XML-based companion file. The CollectControl application also allows users to activate, deactivate, and tune the image stabilization and target tracking functions in the video server and displays boxes around moving targets identified by the target tracker.

New and Modified Software

Several of the assumptions that were made in the design of the agent registry and discovery system proved to be problematic during this development effort. In particular, the concept of a platform as "one or more hosts" was an issue. To minimize the hardware required for the system, it was desired to use one physical server to process the video from all cameras in the system. Since each camera is in a separate location and has an independent pan-tilt unit, the cameras must be associated with different platforms. The agent registry system, however, had no provisions for hosting multiple logical platforms from a single physical host.

To address this shortcoming, modifications were made to the agent registry server as well as the underlying network libraries to support the use of multiple interfaces on a single system. The advanced IP setup options in the Windows operating system were used to add multiple IP addresses to the server. Each individual IP was assigned to a single platform. Multiple instances of the agent registry, as well as all other required software, were then run on the server, each associated with a different platform through environment variable settings. These settings made the server appear as multiple logical platforms to other nodes on the network. All other client and server software required no source code modifications.

MCVideoAgent contained an image processing plug-in that performed moving target identification and tagged the frame with the resulting target data. The algorithm used by this plug-in was simplistic and proved to be ill-suited to real world applications, as it could not correct for camera shake and similar issues. Camera shake can be caused by even low-speed wind when the camera is elevated, and the result is the appearance of movement by all the objects in the image. CI-CB researchers and engineers developed a new moving target tracking and image stabilization plug-in that corrected the image for camera shake before performing target tracking. Further enhancement to this plug-in will enable the target tracker to be used on a moving camera, including cameras mounted to mobile platforms as well as stationary cameras in the process of panning or tilting.

Additional modifications were necessary to support the digital video recorder (DVR) system. By design, MCVideoAgent maintained a single video feed that was sent to all clients. All clients received the feed at the same resolution, frame rate, and image quality. If a client adjusts the feed parameters, all other clients will receive a feed with the new parameters. This design allows for a single compression and video processing thread to be used regardless of the number of connected clients, keeping CPU utilization constant. This design also allows for the use of IP multicast to distribute the video stream on a suitably equipped network.

The DVR, however, requires a constant frame rate for ideal performance. To accommodate this, we modified MCVideoAgent to support multiple feeds with different video parameters from the

same camera. This mode is not used unless the digital video recorder is obtaining video from the server. All clients with the exception of the DVR continue to share a common feed. If connected, the DVR is provided with an independent feed at a constant resolution and frame rate. This approach preserved the scalability of the old design while allowing the DVR to operate at a constant frame rate.

The DVR server reads a list of camera names and associated recording parameters from a configuration file and uses the agent registry server to attempt to locate the corresponding video sources on the network. When a configured video source is found, the DVR server connects to the video source and begins to receive video. The DVR stores incoming video frames, associated posture data, and timestamps to the local hard drive. Video is stored until a specified maximum file size is reached, at which point the oldest video data is overwritten. Optionally, old data can be archived to a second data set at a fractional frame rate before it is overwritten. The DVR uses a custom indexed storage format that facilitates the rapid retrieval of frames of video at a given point in time. Video frames are never recompressed, preserving image quality. The DVR also contains provisions for certain frames to be marked as containing “interesting events”, although this option is currently not fully implemented.

Clients can connect to the DVR server and request video streams or individual frames of video using an XML-based protocol. The server provides clients with a list of available cameras and time ranges. Clients can request a frame or stream starting at any time within the range stored on the server, and can request streams be replayed forwards or backwards in real time or at accelerated or fractional frame rates. Clients can also request to download a block of frames in order to export data to standard file formats.

The DVR server is capable of simultaneously recording an unlimited number of cameras within the constraints of the hardware and network resources. The DVR can be located anywhere on the network. In the persistent surveillance test bed, the DVR is located on the video processing server, but in other deployments the DVR can be located wherever the best connectivity and largest storage capacity is available. Multiple DVR systems can be used in the same network for load balancing or redundancy.

We also developed a client application for the DVR server as part of the persistent surveillance effort which is shown in figure 5. This application presents a video cassette recorder (VCR) - like interface to the operator. The operator can select from any DVR server available on the network and is then presented with a list of cameras the server is configured to record. Upon selection of a camera, live video playback begins and a list of available time ranges is presented to the user. The user can seamlessly switch from viewing live video to recorded video, and can play recorded video forwards or backwards at variable speeds. The user can also rapidly jump to video recorded at any given time as well as “scroll” backwards and forwards in the video stream. The client application also contains functions to export video clips and time-lapse movies as standard AVI files.




Figure 5. DVR client application.

Surveillance Applications

The primary function of the camera surveillance system is to provide image data over a ROI. To this extent, there are two primary measurements used to evaluate the coverage provided: the greatest area of coverage (GAC) and the greatest perimeter coverage (GPC). Both of these measurements have a dependency on the value r , which is the maximum range of the camera¹⁴. However, the maximum range of the camera is not a fixed value. The value of r is dependent on the capabilities of the camera, the intended application of the camera, and the environment. The capabilities of the camera including resolution, maximum zoom, and other features such as infrared vision are considered with the intended application to give a baseline value for the range. The intended use of the camera is also a factor in determining the camera's range. For example, a camera can identify a truck from a greater distance than it can identify a person. In this example, the application of identifying a truck yields a greater range value than if the application were identifying a person because the truck is larger than the person. Conversely, a camera will have a lower range value if it is reading a car's license plate as opposed to identifying a car's color. Here, range is limited because reading the license plate requires a

¹⁴In this paper, a camera's range is defined as the maximum distance image data from the camera is useful for target recognition, classification, and/or identification depending on the desired application.

higher resolution. In addition, environmental factors can affect the range of a camera. A camera can see farther on a clear day than on a foggy night. If the environmental factors change over time, the range of the camera becomes a dynamic value.

GAC is a measurement of the area that the camera system can monitor. Ideally, an individual camera would have a GAC of πr^2 , or the area of a circle centered on the camera location where r is the maximum range of the camera. However, the GAC measurement must also take terrain and other obstructions into account. For example, a tree obscures the area behind its trunk, subtracting from the GAC measurement.

This is shown in figure 6, where the large circle shows the GAC of an ideal system. The darker shading indicates an area that would be included in the ideal GAC, but is subtracted by the obstruction (the larger white circle). In the rooftop camera system, each camera is obstructed by their mounting bracket and the roof itself.




Figure 6. Camera GAC showing an obstruction.

The GAC of the system as a whole includes the GAC of each individual camera. It is imperative to consider the effect of camera placement when creating a system with multiple cameras.

Cameras that are within a distance of $2r$ of each other, where r is the maximum range of a camera, will likely have coverage areas that overlap. This overlap is represented by the darker shaded area in figure 7, which can be viewed by both assets. However, being within a distance of $2r$ of each other does not guarantee overlap. For example, consider a camera that has a GAC of less than πr^2 because of a large tree obstructing some area. If another camera was placed opposite of the tree, the system GAC would now include the area on both sides of the tree. This is an example of how intelligent camera placement can help eliminate obstructed or obscured areas.




Figure 7. Coverage area of two cameras.

The configuration of the individual cameras in this system attempts to maximize the GAC centered on the building with the restriction that the cameras must be placed on the roof. The cameras are placed the maximum distance possible away from each other to minimize the radial overlap (the area covered by πr^2). In addition, the elevation of the cameras increases the area that the cameras can monitor by limiting the obstruction of low elevation obstacles, as demonstrated in figure 8. In the top picture, the initial configuration of a system is shown. In the bottom figure, the effect of raising the camera elevation is demonstrated, as the darker shaded area, which indicated area hidden from the camera, is smaller. Finally, by placing the cameras on the corners of the building, the obstruction caused by the roof is minimized. For example, a camera placed at the center of the roof would not be able to monitor the ground level close to the building.




Figure 8. Elevation view of GAC.

While creating a system with a large GAC may seem like the most effective use of camera assets, this is not always the case. The tradeoff of having a large GAC is the expense of monitoring the area. While a pan/tilt camera may be capable of viewing a large area, it cannot view the entire area at any given time. The camera must instead employ its pan/tilt motion to view a portion of the area it can survey. This process costs time to move the camera as well as processing power to determine where the camera should point. Additionally, a camera may not have the resolution to effectively cover areas on the fringe of its range. For example, a camera may be able to monitor a large tree far into the distance, but may not be able to distinguish a license plate at that distance. Depending upon the requirements of the camera system, this may or may not be sufficient to include this area in the GAC.

The GPC is another measure of how effective the placement of a camera is. In ideal conditions, a camera's GPC is $2\pi r$, the perimeter of a circle centered on the camera, where r is the maximum range of the camera. While the GAC is a measurement of the area monitored, the GPC is a measurement of the border monitored. The area monitored is useful for applications such as tracking a moving object or analyzing terrain data. The border monitored can be used to detect intrusion into and out of an area. In general, a large GAC indicates a large GPC and vice versa, although depending on terrain and camera position, this may not be true.

As was the case with the GAC, the value of the GPC can differ from the estimate of $2\pi r$ due to obstacles and camera overlap. For example, a sharp drop in terrain elevation near the maximum of the camera's range would hide terrain, thus shrinking the border that can be monitored. However, re-placement of the camera, such as raising the elevation, can restore the GPC closer to the ideal $2\pi r$ value. Note that obstacles that reduce the GAC will not necessarily subtract from the GPC, as long as the maximum perimeter can be monitored. For example a small valley will

obscure area and subtract it from the GAC. However, if the camera can still monitor beyond the valley, it will not subtract from the GPC.

Since this system is a test bed, it does not have a single defined task. The system will be used for experimentation to identify or track many different types of targets at different resolutions. Additionally, other assets and behaviors will be added that can affect the abilities of the system. Therefore, an r value for the test bed cannot be defined. Instead, each individual experiment will have its own criteria that determine the r value, and consequently the system GAC and GPC.

Tracking Applications

An important aspect of persistent surveillance is the ability to track a target. In the context of a camera system, maximum time on target (MTT) is the goal of tracking a target for as long as possible, or until it is no longer of interest. The MTT of a system is not a hard number, as the ability to track a target depends on the target's velocity, acceleration, and heading. Instead, the MTT is a goal when designing a surveillance system. MTT can be maximized by reducing or eliminating "blind spots", or areas encompassed by but not included in the GAC. Such "blind spots" can include the area directly behind a tree or hill. Such areas allow targets to hide and avoid the surveillance cameras. A target can also avoid surveillance by simply moving outside of a camera's GAC. In an ideal surveillance system, a target will only exit the GAC if it enters an area outside of the ROI, so that it is no longer considered to be a target. A third scenario in which a target could evade the surveillance system is when the target moves at a rate that the surveillance devices cannot track. While the target is inside of the GAC, the camera must pan and tilt to keep the target within the field of view. A target could potentially move so quickly that the camera cannot pan and tilt at a sufficient speed to keep the target within view. Alternatively, the target could be moving faster than the operator or an automated camera controller can process the camera movements and make cueing decisions. To help minimize these potential problems, the system can employ maximum eyes on target (MET).

MET is the goal of viewing a target with multiple cameras. Using multiple cameras to track a target yields several advantages. One advantage is to limit the "blind spots" that occur from areas that are obscured from one or more cameras. When the target is viewed from several angles, there are fewer "holes" in which the target can hide. Another advantage of having a large MET is that faster moving targets can often still be tracked. If the target moves fast enough to evade one camera, it is possible that the other cameras also tracking the target could still keep the target in view.

The most direct application of MET is the ability to "hand off" the target from one camera to another. In a multiple camera system, each individual camera will have its own GAC that it contributes to the system's overall GAC. A target can leave the individual camera's GAC while

staying in the system's GAC. If only a single camera tracks the target, when the target leaves the camera's GAC, there may be a delay while the next camera pans and tilts to bring the target into view. With MET, this "hand off" will occur while keeping the target in view.

The drawback of using MET is the allocation of resources to track a single target. The greater the number of cameras tracking a single target, the larger the area that is left unmonitored by these occupied cameras. In addition, processing power or operator attention must be allocated to each camera so that it can track the target. This could strain the computing power of the system if all camera tracking is handled by a single source.

The ALC test bed has overlap between each camera's GAC, but because of the obstruction caused by the roof itself, the roof can only have a maximum of three cameras on any given target. The exceptions to this rule are if the target is on the roof itself, or if the target has a greater elevation than the roof, such as a helicopter.

Heterogeneous Asset Collaboration

The cameras used in the ALC rooftop system are tied to a central server, where an individual camera can be cued to pan, tilt, and zoom through either software behavior or manual control. The goals for improving the system include integration with additional sensors, collaboration with mobile assets, and improvements to tactical behaviors.

The current system only provides visual data from the four rooftop cameras. The video feeds are processed at a central server, where they can be sent to individual clients. The cameras can be cued to manually pan, tilt, and zoom by any of these clients. In addition, a software package can be employed to track movement. Planned future software development will also allow the cameras to be cued to follow movement. The cameras can also be cued according to data obtained through an acoustic array meant to locate the direction of sniper or mortar fire. Future improvements will integrate additional sensors. For example, an unmanned ground sensor (UGS) can cue the system to point any available camera to focus on an area upon detecting movement. Also, a Soldier with a Global Positioning System (GPS) device and appropriate communications equipment can cue a camera to fixate on himself if he has found a target of interest.

Integration of mobile assets can greatly increase the capabilities of the system. For example, if a camera is tracking a target that moves behind an obstacle into a blind spot, the system can cue a mobile asset to move into position to bring the target back into view. This is especially useful for terrain that has numerous blind spots. In the future, assets could be moved either by manual operator control or by using autonomous navigation technologies as they become available.

Improved tactical behaviors can allow the system to provide more effective surveillance of a given area. The goal of these improved behaviors would be to have the system coordinate all of its assets to identify and track targets effectively. An example of such a scenario would be to identify a target with an infrared tripwire, track the target through movement of a stationary rooftop camera, cue a mobile asset to follow the target, and handoff the target to another camera as it moves into another region.

Conclusion and Future Work

The prototype system successfully provided a test bed for the development, implementation, and testing of persistent surveillance algorithms. The camera locations provide surveillance coverage of all sides of the building, and the elevation of the cameras minimizes blind spots within the ROI. As algorithms are developed and implemented, the system can be easily expanded by adding other assets, such as the mobile assets previously described, and the ROI can be increased to provide a reasonable test bed for the expanded system.

Several enhancements to the system can be made to make it more robust. The current system design requires an AC power source for each camera. While this requirement is satisfactory for a roof-top mounted system, it severely inhibits the possible scenarios in which the cameras can be employed. The system could be modified to draw power from a battery. This would allow for the cameras to be mobile or field deployable. However, the change introduces limitations due to limited battery life.

Another potential enhancement is to couple a battery with a solar panel. The addition of this self-replenishing power system increases the system's ability to function without an outside power source for a prolonged period. The roof-top applications that the system was originally designed for would allow for ample sunlight for the panel. A drawback of such an improvement would be the high cost of purchasing and installing solar panels. In addition, the system would be larger and require more planned placement to get sun exposure on the solar panels. Further investigation is required to determine the size of solar panel and battery needed to meet the system's power demands.

Another limitation of the current system is the fiber optic cable used for communications between the cameras and the central server. Once again, this requirement was satisfactory for a rooftop system that can provide existing infrastructure support. To make the system easier to deploy, the fiber optic cable could be replaced by a wireless network. Other current areas of research for CI-CB include mobile ad-hoc networks (MANET). This technology could potentially be leveraged to make the system more robust. Adding wireless capabilities would increase the system's ability to be rapidly relocated or field deployed. The combinational

improvements of solar power and wireless capabilities would allow for remote deployment of the camera systems for long periods of time.

Future research efforts in software design will also result in improved capabilities within the persistent surveillance system. ARL is collaborating with the Institute for Human and Machine Cognition, part of the Florida University System, to develop a decentralized, peer-to-peer discovery architecture to replace the existing hierarchical agent registry server. This new discovery system uses IP multicast or a data-aware network to distribute service information and requests among nodes in the network. The new architecture will remove the static configuration requirement necessary to initialize the agent registry server. Additionally, the new system will provide better performance in unreliable network environments, such as MANET's, where end-to-end connectivity is not guaranteed. When integrated into the persistent surveillance system, this new discovery architecture will allow for easier and quicker addition of new assets to the system. It will also enable the surveillance system to discover and make use of resources on mobile assets that may only be connected to the network part of the time.

INTENTIONALLY LEFT BLANK.

DISTRIBUTION LIST

ADMNSTR
DEFNS TECHL INFO CTR
ATTN DTIC-OCP (ELECTRONIC COPY)
8725 JOHN J KINGMAN RD STE 0944
FT BELVOIR VA 22060-6218

DARPA
ATTN IXO S WELBY
3701 N FAIRFAX DR
ARLINGTON VA 22203-1714

OFC OF THE SECY OF DEFNS
ATTN ODDRE (R&AT)
THE PENTAGON
WASHINGTON DC 20301-3080

US ARMY TRADOC
BATTLE LAB INTEGRATION & TECHL
DIRCTRT
ATTN ATCD-B
10 WHISTLER LANE
FT MONROE VA 23651-5850

PM-FCS (BCT) TECHNOLOGIES
ATTN M/S CV-52 S YOUNG (2 COPIES)
7990 SCIENCE APPLICATIONS CT
VIENNA VA 22182

SMC/GPA
2420 VELA WAY STE 1866
EL SEGUNDO CA 90245-4659

COMMANDING GENERAL
US ARMY AVN & MIS CMND
ATTN AMSAM-RD W C MCCORKLE
REDSTONE ARSENAL AL 35898-5000

US ARMY INFO SYS ENGRG CMND
ATTN AMSEL-IE-TD F JENIA
FT HUACHUCA AZ 85613-5300

US GOVERNMENT PRINT OFF
DEPOSITORY RECEIVING SECTION
ATTN MAIL STOP IDAD J TATE
732 NORTH CAPITOL ST., NW
WASHINGTON DC 20402

US ARMY RSRCH LAB
ATTN AMSRD-ARL-CI-OK-TP TECHL
LIB
T LANDFRIED (2 COPIES)
BLDG 4600
ABERDEEN PROVING GROUND MD
21005-5066

US ARMY RSRCH LAB
ATTN AMSRD-ARL-CI J GOWENS
ATTN AMSRD-ARL-CI-C B BROOME
ATTN AMSRD-ARL-CI-CB B O'BRIEN
(3 COPIES)
ATTN AMSRD-ARL-CI-CB C
WINSLOW
ATTN AMSRD-ARL-CI-CB D BARAN
(3 COPIES)
ATTN AMSRD-ARL-CI-CB H NGUYEN
ATTN AMSRD-ARL-CI-CB J D
WALRATH
ATTN AMSRD-ARL-CI-CB J KOVACH
(3 COPIES)
ATTN AMSRD-ARL-CI-CB L
TOKARCIK
ATTN AMSRD-ARL-CI-CB N FUNG
(3 COPIES)
ATTN AMSRD-ARL-CI-CB P DAVID
ATTN AMSRD-ARL-CI-CB R
GREGORY
ATTN AMSRD-ARL-CI-CB R WINKLER
ATTN AMSRD-ARL-CI-CB S
BRAGONIER
ATTN AMSRD-ARL-CI-CB S CHOY
ATTN AMSRD-ARL-CI-CB S HO

US ARMY RSRCH LAB (CONT'D)
ATTN AMSRD-ARL-CI-CB T GREGORY
ATTN AMSRD-ARL-CI-CB T ROSE
ATTN AMSRD-ARL-CI-OK-T TECHL
PUB (2 COPIES)
ATTN AMSRD-ARL-CI-OK-TL TECHL
LIB (2 COPIES)
ATTN AMSRD-ARL-D J M MILLER
ATTN IMNE-ALC-IMS MAIL &
RECORDS MGMT
ADELPHI MD 20783-1197